USER LEVEL CONTROL OF POWER MANAGEMENT POLICIES

BACKGROUND

[0001] This application is a continuation of U.S. patent application Ser. No. 14/855,553, filed Sep. 16, 2015, which is a continuation of U.S. patent application Ser. No. 13/782,473, filed Mar. 1, 2013, now U.S. Pat. No. 9,170,624, issued Oct. 27, 2015, which is a continuation of U.S. patent application Ser. No. 13/326,586, filed Dec. 15, 2011, now U.S. Pat. No. 9,098,261, issued Aug. 4, 2015, the content of which is hereby incorporated by reference.

Background

[0002] Advances in semiconductor processing and logic design have permitted an increase in the amount of logic that may be present on integrated circuit devices. As a result, computer system configurations have evolved from a single or multiple integrated circuits in a system to multiple hardware threads, multiple cores, multiple devices, and/or complete systems on individual integrated circuits. Additionally, as the density of integrated circuits has grown, the power requirements for computing systems (from embedded systems to servers) have also escalated. Furthermore, software inefficiencies, and its requirements of hardware, have also caused an increase in computing device energy consumption. In fact, some studies indicate that computing devices consume a sizeable percentage of the entire electricity supply for a country, such as the United States of America. As a result, there is a vital need for energy efficiency and conservation associated with integrated circuits. These needs will increase as servers, desktop computers, notebooks, ultrabooks, tablets, mobile phones, processors, embedded systems, etc. become even more prevalent (from inclusion in the typical computer, automobiles, and televisions to biotechnology).

[0003] In many computing environments, it is an established fact that for much of the time, computing systems such as servers are operating well below their peak performance level. During these periods of low utilization the focus is on saving as much power as possible in order to reduce the energy costs. Power management technologies can deliver significant power savings during periods of low utilization. However any power management technology involves a power/performance tradeoff.

[0004] Due to increasing integration, many processors can include power management technologies which can control up ⅔ rds of total platform power. In many cases these technologies are controlled by a power control unit (PCU) in the processor. Each power management feature is specifically tuned in design to achieve an optimal power/performance tradeoff. At the time of tuning, there is little knowledge of the actual workload and usage pattern for the system in the field. Given this lack of knowledge, the tuning process is conservative and is necessarily biased towards losing as little performance as possible. This approach prevents significant power savings for an end user who is willing to tolerate more performance loss in return for power savings.

[0005] Thus typically power management features are statically tuned to tolerate very little performance loss. This results in several negative downsides. First, at low utilizations where an end user can tolerate high performance loss, available power savings are not realized. Second, an end user typically has no choice regarding power/performance tradeoffs, other than default profiles provided by an operating system (OS). Given the complexity involved in tuning power management features, end users rarely venture into tuning individual features for their target usage, and thus the potential benefit of the features are often not realized.

BRIEF DESCRIPTION OF THE DRAWINGS

[0006] FIG. **1** is a block diagram of an architecture of a tuning circuit for a tunable power performance loadline technique in accordance with one embodiment of the present invention.

[0007] FIG. **2** is a block diagram of a tuning table in accordance with an embodiment of the present invention.

[0008] FIG. **3** is a flow diagram of a method in accordance with an embodiment of the present invention.

[0009] FIG. **4** is a block diagram of a processor in accordance with an embodiment of the present invention.

[0010] FIG. **5** is a block diagram of a multi-domain processor in accordance with another embodiment of the present invention.

[0011] FIG. **6** is a block diagram of a system in accordance with an embodiment of the present invention.

[0012] FIG. **7** is a block diagram of a multiprocessor system with a point-to-point (PtP) interconnect in accordance with one embodiment of the present invention.

[0013] FIG. **8** is a block diagram of a partially connected quad processor system in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION

[0014] Embodiments provide a so-called energy performance bias (EPB) as an architectural feature. Control of this parameter allows for a simple high level input from an end user to indicate a power/performance tradeoff preference from the end user. This input can be used to provide multiple tuning levels with different points of power and performance tradeoff. By associating this energy performance bias with direct user input, embodiments enable the end user to directly control power/performance tradeoff in a simple manner. As used herein the terms "end user" or "user" are comprehended to include computer users of varying degrees, including technical and non-technical users, information technology (IT) personnel, data center personnel and so forth.

[0015] Thus instead of providing complete tuning flexibility for each power management technology and allowing an end user to tune each feature, a single input can be provided by the user to control these different features. The EPB value may thus correspond to a single input value to control a plurality of power management features. Furthermore, understand that the provision of the EPB value can be from a variety of external entities including but not limited to an operating system (OS), a basic input/output system (BIOS), an external embedded controller of a platform such as a baseboard management controller (BMC), a data center central management software and communicated via a network and a node manager device or so forth to a platform, among others, automatically or via a user. And in some embodiments, the end user may be prevented from such individual control. As such, the inherent difficulty in exposing all of a large number of power management features to the end user can be avoided, particularly as most end users